



Adaptive Prompting for Clinical Text Summarization: A Few-Shot Learning Approach Using Generative Large Language Models

D. Vithanage^{1}*

¹ Department of Information and Communication Technology, Faculty of Technology, University of Ruhuna, Matara, Sri Lanka

**Corresponding author: dinithi@ictec.ruh.ac.lk*

Clinical text summarization is a critical yet challenging task due to the unstructured and heterogeneous nature of electronic health records (EHRs). Discharge summaries, in particular, often contain lengthy and complex narratives, making it difficult to efficiently extract key clinical insights. Recent advancements in generative large language models (LLMs) have shown great promise in automating summarization tasks; however, their optimal application in clinical contexts remains underexplored, representing a significant research gap. To address this gap, the present study investigates an adaptive prompting approach for discharge summary summarization using different few-shot learning settings, ranging from one-shot to six-shot, with several state-of-the-art generative LLMs, including GPT-5, LLaMA, Mistral, Gemini, and T5. The performance of these models was evaluated using multiple metrics—such as F1 score, ROUGE, and BERTScore—alongside complementary measures like BLEU, BLEURT, and METEOR, to ensure a comprehensive assessment. The experimental results revealed that GPT-5 achieved superior performance across all evaluation metrics in the five-shot configuration, reaching an average score of 99%. This outcome underscores the strong effectiveness of five-shot learning for clinical summarization tasks. These findings provide valuable insights for healthcare researchers, informatics practitioners, and policymakers aiming to integrate LLMs into clinical workflows to support efficient information retrieval and informed clinical decision-making.

Keywords: Large language models, Clinical text summarization, Electronic health records